Journal of Intelligent Information Systems
https://doi.org/10.1007/s10844-023-00801-4

RESEARCH

Session-based recommendation with temporal dynamics for large volunteer networks

Taurai Muvunza[1] · Yang Li[1]

Received: 7 February 2023 / Revised: 30 May 2023 / Accepted: 31 May 2023

Abstract
In large online volunteer systems, inefficiency and low volunteer retention are existing challenges that compromise the success of online communities particularly given the uncertainty in volunteer participation behavior. A strategy that matches volunteers to a host of fields will alleviate these challenges, yet creating an all-in-one volunteer recommendation system is an unexplored but promising area. We propose VolRec, a session-based recommendation for large volunteer networks that employs temporal dynamics to capture uncertainty caused by the changing structure of volunteers' participation behaviour. To optimize the recommendations, we construct a probabilistic volunteer network graph that denotes co-participation in an activity. We then model individual and inferred neighbours' preferences as dynamic and context-aware sessions. VolRec can be adapted to recommend volunteers to organizers, tasks, groups and communities, creating a comprehensive and efficient recommendation system. Experiments using Pioneers data, a mobile based app launched in the wake of Covid-19 to mobilize volunteers and record their participation activities demonstrate the efficacy of this approach.

Keywords Session-based recommendation · Volunteer · Graph neural networks · Temporal dynamics

1 Introduction

In recent years particularly in the wake of Covid-19 pandemic, governments turned to online volunteer platforms to coordinate volunteer efforts. As countries grappled with the overwhelming demand for help, crowdsourcing volunteer platforms became pivotal in mobilizing volunteers to effectively control the pandemic (Marston et al., 2020). However, when the crowdsourcing platform is not properly designed, volunteers spend a lot of time searching

✉ Yang Li
yangli@sz.tsinghua.edu.cn

Taurai Muvunza
gnf18@mails.tsinghua.edu.cn

1 Tsinghua-Berkeley Shenzhen Institute, Shenzhen International Graduate School, Tsinghua University, University Town, Shenzhen 518055, Guangdong, China

for active organizers, tasks and communities to participate in. This negative experience leads to dissatisfaction and results in exodus of both organizers and volunteers from the platform, plummet volunteers' retention rate and trigger the ultimate collapse of the system, (Locke et al., 2003). To improve the efficiency of crowdsourcing systems, it is crucial to understand volunteer participation behavior and learn their representation to recommend volunteers to a host of possible fields. Our study is motivated by the growing adoption of crowd source volunteer platforms, where individual volunteers collaborate to achieve rapid mobilisation toward emergent community demands.

Yet there has been fundamental challenges in designing optimal recommendation systems due to uncertainty in user behaviors. Given user-item interaction matrix, existing collaborative filtering methods such as ItemKNN and Matrix factorization use this matrix to learn embeddings and rely on similarity to compute predictions. This approach is problematic since it does not accommodate the sequential order of users' consumption of items as well as their social network relations which could be very informative in predicting the next item. A recent study by Lei et al. (2022) utilizes sequential user-item interactions but ignores social relationships when making recommendations. These limitations create sub-optimal recommendations and this concern motivates our work.

To address the above limitations, we utilize temporal-dynamic information to create sequences of volunteer participation. To guarantee optimality, we segment these sequences into sessions using an entropy based method. Additionally, we incorporate social relationships by creating a probabilistic volunteer network graph based on users' co-participation history. To capture the dynamic relationships in the volunteer network, we employ Graph Neural Networks (GNNs) and propagate the influence of each neighbour for a target volunteer agent using an attention mechanism. As we will show, dynamic volunteers' task participation can be modeled as session embeddings, and including volunteer network information significantly improves the recommendation accuracy by $\sim 26\%$ as partly displayed in Fig. 1.

We study online crowdsourcing volunteer systems because they are less explored, yet seeing a rebound in use since Covid-19 pandemic (Marston et al., 2020). For instance, when

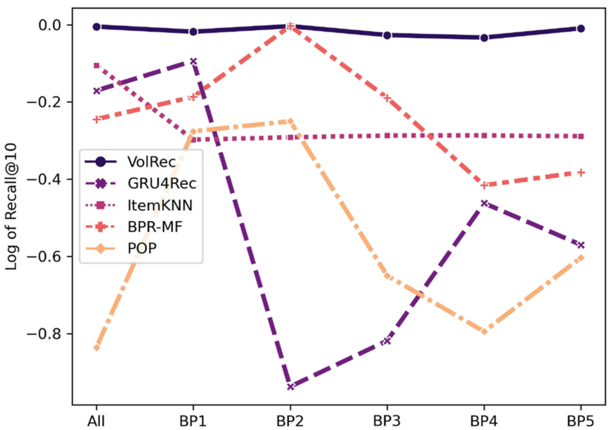

Fig. 1 Visualization of results for VolRec compared against existing baselines. We evaluate the models on volunteer-organizer recommendation using the log of Recall@10 and compare their performance under different scenarios. Our model outperforms the baselines both when all data is used, and on structural break point intervals, BP1-5, computed at district level

an organizer-volunteer recommendation system is set up effectively, on one hand volunteers spend less time finding tasks, and on the other hand, organizers are assured to get the right volunteers for their tasks. From a policy perspective, the ability to predict volunteer participation particularly as sessions is vital for local authorities and community leaders to determine and allocate the precise number of resources required for each task. Using data collected from a large-scale mobile volunteer platform (Zhang et al., 2022) in Shenzhen, China, we learn the volunteer-organizer relationships in a decentralized setting where a volunteer can also be an organizer.

This study makes the following contributions: 1) We proposed VolRec, a sequential recommendation algorithm as a strategy to deal with uncertainty and inefficiency in large crowdsourcing volunteer platforms. 2) To optimize the recommendations, we construct a probabilistic volunteer network graph from users' participation history and further segment volunteer interactions with organizers, tasks and communities into context aware sessions, which consist of ordered sequences that are learned and propagated using an attention mechanism. 3) We adopted GNNs to capture high order structural data in large volunteer networks. 4) We conduct extensive experiments using real world volunteer data. Experimental results demonstrate the effectiveness of our session-based recommendation with temporal dynamics over existing baselines. Moreover, this study is the first to propose a recommendation system as a strategy to improve efficiency of large volunteer networks.

The rest of the paper is organized as follows: Section 2 recaps related work in recommender systems. Section 3 describes the methodology for VolRec while Section 4 discusses experiments and subsequent results. We conclude in Section 5.

2 Related work

2.1 Recommendation systems

Recommendation systems have important applications in a range of areas including social networks (Wang et al., 2014); online purchases (Linden et al., 2003); movies, online bookmarking and restaurants (Lei et al., 2022; He et al., 2022; Song et al., 2019); music streaming (Choi et al., 2022; Guo et al., 2019); web browsing (Xu et al., 2018) and news (Liu et al., 2010).

A significant body of scholarship has proposed several recommendation systems in literature. Generally, recommender systems fall into three categories namely shallow, neural and GNNs models, (Gao et al., 2022). Shallow models include Matrix factorization algorithms (Koren et al., 2009; Rendle et al., 2010), Item-based collaborative filtering methods (Linden et al., 2003; Sarwar et al., 2001) and Personal Ranking methods (Rendle et al., 2012). While this class of models are easy to compute, they have severe shortcomings. For instance, ItemKNN treats user-item interactions as static over time and does not consider complex user behaviors. More importantly, Matrix Factorization (MF) methods model user and item features by computing the inner product between the latent features of users and items. Neural network models such as Neural Collaborative Filtering (NCF) (He et al., 2017) improved the shallow models by replacing the inner product with multi-layer perceptrons (MLP) that learn the user-item interaction function. Neural models achieved superior performance due to the use of deeper layers and their ability to learn non-linear relationships. However, these methods are limited since they do not capture the high-order structural information in data, (Gao et al., 2022). For instance, NFC recommends items by optimizing the prediction of items that the user has interacted with, while ignoring neighbourhood information.

2.2 Session-based recommendation

Previous studies have introduced a number of cutting edge technologies in sequential recommender systems. Session-based recommendation is a form of sequential modeling that seeks to model a user's temporal behaviour by capturing the user's successive sequential consumption of a basket of items. The differences between session-based recommendation and sequential recommendation are formally outlined in a survey by Wang et al. (2022). Prior research on sequential recommendation has used powerful methods based on Recurrent Neural Networks (RNN) and its variants, Long Short-Term Memory (LSTM) (An et al., 2019; Li et al., 2017) and Gated Recurrent Unit (GRU) (Hidasi et al., 2016; Donkers et al., 2017). Besides, other studies have used Markov Decision Processes (MDP) (Shani et al., 2005) and random walk (Choi et al., 2022) to model the sequential order of user-item interactions.

Additionally, state-of-the-art models such as NARM (Li et al., 2017) and STAMP (Liu et al., 2018) have also achieved significant results. NARM utilizes global and local encoders to capture user's sequential behaviour while STAMP applies a simple MLP and attentive net to model general and current interests of the user. SR-GNN (Wu et al., 2019) models session sequences, represented by a combination of global and local preferences as session graphs and use GNN to capture the item transitions. Recommender systems have also been adapted to specific scenarios. For instance, Guo et al. (2019) proposed SSRM, which integrates MF techniques with session-based recommendation to model user behaviour in online music streaming. Despite demonstrating high levels of performance, these models neglect latent relationships between users. Incorporating implicit social network data in user modelling is still an emerging area of interest in recommender systems (Song et al., 2021). A recent survey by Gao et al. (2023) provides an overview of deep learning methods tailored to various contexts including social recommendation. Several conventional methods such as Popularity Predictor (POP), ItemKNN (Linden et al., 2003) and BPR-MF (Rendle et al., 2012) are established baselines with cutting-edge performance in recommender systems.

2.3 Recommendation on social networks

Studies have shown that people's networks are homogenous, and a user is similar to or influenced by the people around him/her, (McPherson et al., 2001). In addition, recent research in recommender systems has exploited social relations and dependency connections and recorded significant improvements, (Fan et al., 2019). While previous studies suggest that incorporating social network information enhances learning and representation of a user's latent space, little work has been done to derive and extract meaningful neighbourhood features where they are not initially given. On another front, Eagle et al. (2008) explored observable behaviors such as proximity between people to infer friendship and found a correlation between physical proximity and the probability of friendship.

2.4 Graph attention networks

GNNs have achieved state-of-the art in many applications mainly because of their graph structure that allows the propagation of high-order neighbourhood information. Several variations of GNNs have been proposed. For instance, Kipf and Welling (2016) introduced Graph-Convolutional Networks (GCN) for semi-supervised graph classification where node representations are updated by aggregating the current node's representations and that of its neighbours. Consequently, network architectures based on GCN, for example, GraphSage (Hamilton et al., 2017) attributes a static weight to each neighbour when updating node representations. Although GCNs captures dependencies among nodes, the influence of each neighbour on the current node is not always identical and deterministic. To address this challenge, Veličković et al. (2017) proposed Graph Attention Networks, an architecture that adopts an attention mechanism to weigh the influence of neighbours differently.

2.5 Online volunteer systems

Online crowdsourcing volunteer systems are less explored in literature. The most recent work in this field is by Zhang et al. (2022) who investigated how self-organizing emerged during the Covid-19 pandemic response and its effectiveness in crowdsourcing volunteer organizations. Our study builds on the inspiring work of Song et al. (2019) and Wu et al. (2019) on sequential recommendation. The authors extract features from predefined social relationships without employing techniques to optimize the network for better performance. With the complexity of online crowdsourcing data, it is increasingly important to learn the latent structure of the relationships within data when building online recommendation systems.

3 Methodology

In this section we begin by introducing the probabilistic volunteer network graph and show the subsequent neighboring effect of volunteers. Next, we present the architecture for our proposed framework. Last, we describe the modeling of volunteers and their neighbours' long and short term preferences and the propagation strategy used in learning these features in the neural network.

3.1 Volunteer network

We construct the volunteer network graph, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where nodes $\mathcal{V}$ indicate target volunteers and their neighbours; and edges $\mathcal{E}$ correspond to their friendship. We build the graph in two steps. First, we compute a 'bag of activity' $\mathcal{Q}(v, o)$ denoting volunteer-organizer interaction to record the frequencies of volunteers' v participation in organizers tasks. $\mathcal{Q}(v, o)$ is a $V \times \mathcal{O}$ matrix, where $\mathcal{O}$ is the total number of organizers. We further express $\mathcal{Q}(v, o)$ as a probability matrix consisting of $P(v \mid o)$. In the second step, for each volunteer $v \in V$, we define cosine similarity as follows:

$$S(q_{v_i}, q_{v_j}) = \frac{\sum_{i=1}^{n} q_{v_i} q_{v_j}}{\sqrt{\sum_{i=1}^{n} q_{v_i}^2}\sqrt{\sum_{i=1}^{n} q_{v_j}^2}}, q_{v_i} \neq q_{v_j} \tag{1}$$

$S(q_{v_i}, q_{v_j})$ represents the similarity between volunteer v_i and v_j's task participation behaviour, where q_{v_i} and q_{v_j} are denoted as $\mathcal{Q}(v_i, .)$ and $\mathcal{Q}(v_j, .)$, respectively. We build the adjacency matrix by sorting all nodes v with respect to $S(q_{v_i}, q_{v_j})$ such that:

$$A = \begin{cases} 1 & \text{if } S(q_{v_i}, q_{v_j}) >= S(q_{v_i}, q_{v_k}) \\ 0 & \text{otherwise.} \end{cases}$$

where q_{v_k} is v_k's bag of activity embedding and v_k is the $k-$th most similar node measured by S with respect to v_i. For example, we fix the number of v_i's neighbours empirically to $n = 10$. The adjacency matrix is a $V \times 10$ matrix consisting of V volunteers and their corresponding top 10 neighbours.

The findings of Eagle et al. (2008) and Wang et al. (2014) on friend recommendation inspired our approach to use volunteers' task participation history to infer social relationships and the use of cosine similarity allows us to capture the latent relationships such as proximity of volunteers and shared task preferences. The probabilistic rendering of volunteer network yields embeddings that are more representative of our volunteers' task participation behaviour given an organizer.

Each node in the graph utilizes its respective volunteer's representations as features. We utilize this graph to propagate the influence of each neighbour in the network towards a target volunteer. We note that this approach is a special case as it differs from the mainstream session-based recommendation systems that base their recommendations solely on pre-existing social networks (Song et al., 2019). To our knowledge, session-based recommendation previously has not been applied in this fashion to online volunteer platforms although it has seen applications in social networks (Wu et al., 2019; Fan et al., 2019). In this study, we model volunteer participation as a sequence of organiser tasks 'consumed' in a session and predict the list of ranked top K organisers. We define the neighbouring effect on a target volunteer's choice and our motivation for building the volunteer network graph as depicted in Fig. 2.

In order to fully utilize a volunteer's information to improve the recommendation, we incorporate volunteer network data into the recommendation. For instance, Fig. 2 illustrates that in Session 1, Alice participates in two tasks under two different organizers. However, Alice's next choice may be influenced by her neighbours' interests, Wang, Pam and Bob. In Session 2, Alice's preferences on task participation changes, and so does her neighbourhood's preferences. Suppose that Alice's most recent participation is organized by Susan, (who leads Covid-19 related tasks) and that Wang is the only one with long term interest of working with Susan. Since none of the other neighbours has participated in Susan's tasks, it would be more plausible for Alice to be influenced more by Wang in Session 2. Therefore, Alice's choices are

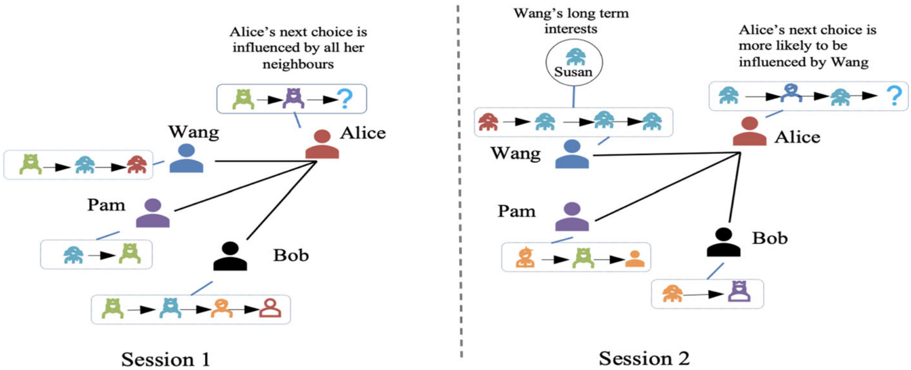

Fig. 2 Illustration of the neighboring effect on a target volunteer's organizer choice. The icons in each box correspond to the sequence of organizers for each task that the volunteer has taken part in. The left side shows that Alice's task participation preferences are influenced by the behavior of Wang, Pam and Bob in Session 1. On the right side in Session 2, neighbours' and target volunteers' preferences change, and Alice's choices are most likely to be influenced by Wang's long term interests but barely influenced by Pam and Bob

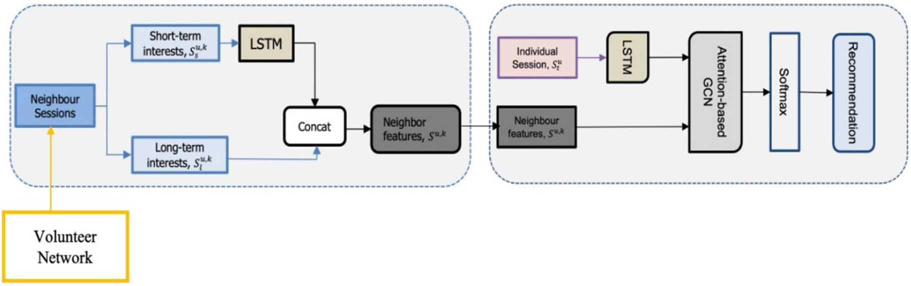

Fig. 3 The architecture for VolRec. We construct the volunteer network graph and extract neighbourhood features by modeling their interests separately as long and short-term sessions. The final node representation is obtained by concatenating the neighbourhood and target node's features using a fully connected network

influenced by both the long and short-term interests of her neighbours. We account for these changes in dynamic volunteer interests by modeling individual and neighbours' preferences separately.

3.2 Overview of the framework

Figure 3 shows the architecture of our proposed approach. One of our main contributions is the construction of a probabilistic volunteer network to determine the neighbours. We then separate the neighbour's preferences into short and long term sessions. The short-term interests are modeled using LSTM, and we obtain the neighbour's node features for each layer by concatenating the preferences using a non-linear transformation. Individual volunteer interests are represented by short-term sessions that are modeled using LSTM to obtain target volunteer features, which are then passed together with neighbours' features as inputs to the attention-based GCN. In the following section, we formally introduce the approach used to model individual and neighbours' interests.

3.3 Individual interests

We define a session as :

$$S_T^u = \{S_1^u, \ldots, S_T^u\} \tag{2}$$

and each session

$$S_t^u = \{o_{t,1}^u, \ldots, o_{t,N_{u,t}}^u\} \tag{3}$$

is a sequence of volunteer behaviors where N_u is the total number of organizers in that session. Figure 4 illustrates volunteer behaviour as sessions. In essence, a session S for volunteer u is a set of organizers o, that the volunteer has participated under. $o_{t,p}^u$ denotes the pth organizer of volunteer u at time t, represented as follows:

$$o_{t,p}^u \triangleq \begin{cases} 1 & \text{if } \mathcal{Q}(v, o) > 0 \text{ at time } t \\ 0 & \text{otherwise.} \end{cases}$$

In addition, we separate sessions into long and short term sessions.

The goal of session-based recommendation is to predict the next session given a set of Alice's current session and social network. We express Alice's organizer-interactions as

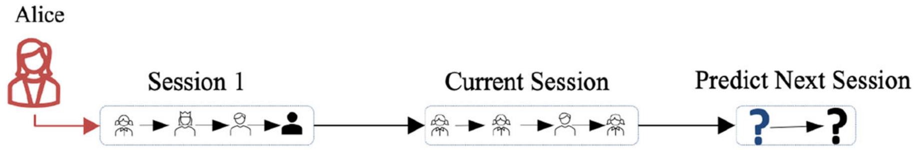

Fig. 4 Each session is a sequence of volunteer participation behavior

short-term dynamic interests modeled using RNN, represented by her latest or current session. However, a volunteer's long term preferences are expressed as the time insensitive average interests that include all historical organizer interactions. RNN infers a user's current session: $S^u_{T+1} = \{o^u_{T+1,1}, \ldots, o^u_{T+1,n}\}$, by recursively combining representations of all previous tokens with the latest token:

$$h_n = f(o^u_{T+1,n}, h_{n-1}) \tag{4}$$

where h_n represents the user's interests and $f(.,.)$ is a non-linear function (LSTM). The initial input features are represented by h_o.

3.4 Neighbours interests

We categorize neighbours' interests into long and short term interests. We model short term interests of a neighbour using RNN. Furthermore, given a target volunteer u with a current session $S^u_{T+1} = \{o^u_{T+1}, \ldots, o^u_{T+1,n}\}$, the neighbour's short term session is represented as the most recent session just before time $(T+1)$. Precisely, we denote the short-term interests of neighbour k for a target volunteer u as:

$$S^{u,k}_s = \{o^{u,k}_{T,1}, \ldots, o^{u,k}_{T,N_{k,T}}\} \tag{5}$$

where the subscript s represents short-term interest while each token corresponds to the neighbour's task participation, and finally N_k is the total number of tasks in the session. Meanwhile, long-term neighbours' preferences are represented as time insensitive average embeddings since they reflect average interests. $S^{u,k}_l$ is the neighbour's long-term preference which corresponds to the k-th row of the user embedding matrix $\boldsymbol{M}_u$. We formally define neighbour's long term preferences as follows:

$$S^{u,k}_l = \boldsymbol{M}_u(k), \tag{6}$$

We obtain the final representation of a neighbour by concatenating the short and long-term preferences using a non-linear transformation

$$S^{u,k} = \mathrm{ReLU}(\boldsymbol{W}_1[S^{u,k}_s, S^{u,k}_l]) \tag{7}$$

where $\boldsymbol{W}_1$ is the transformation matrix and $\mathrm{ReLU}(x) = \max(0, 1)$.

3.5 Graph attention network

Graph Attention Networks (GAT) have proven to have several advantages over other network architectures. For instance, while GraphSage (Hamilton et al., 2017) assumes neighbours' contributions to target node are identical and deterministic, GAT adopts attention mechanism to learn the relative weights between connected nodes (Wu et al., 2019). We adopt GAT and features of each node are propagated using an attention mechanism as shown in Fig. 5.

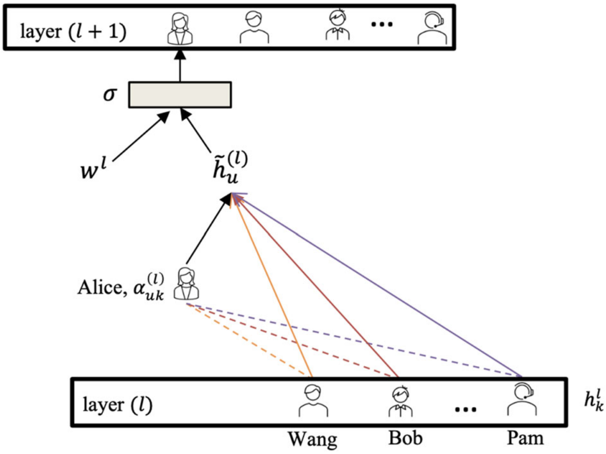

Fig. 5 Single convolutional layer showing the propagation of volunteer network information using an attention mechanism

Figure 5 shows that a target user's behavior depends on his/her individual preference and social influence from neighbours. Therefore, a unified representation of a volunteer's preference is obtained by using an attention mechanism to determine the weights. The dotted lines denote the level of influence of a neighbour towards a given target node, in this case, Alice. Formally,

$$\alpha_{uk}^{(l)} = \frac{\exp\big(f(h_u^{(l)}, h_k^{(l)}\big)}{\sum_{j \in N(u) \cup u} \exp\big(f(h_u^{(l)}, h_j^{(l)})\big)} \tag{8}$$

where $f(h_u^{(l)}, h_k^{(l)}) = h_u^{(l)^T} h_k^{(l)}$ is the similarity between target volunteer and all aggregated neighbours' preferences and $\alpha_{uk}^{(l)}$ represents the weight of neighbour k on target volunteer u. Moreover, to preserve a volunteer's revealed interest, a self-connecting edge, $\alpha_u^{(l)}$, is also included. The neighbours influence and the target volunteer's self-connection edge are combined as follows:

$$\tilde{h}_u^l = \sum_{k \in N(u) \cup u)} \alpha_{uk}^{(l)} h_k^{(l)} \tag{9}$$

where $\tilde{h}_u^l$ is the the mixture of a target volunteer's neighbours at layer l. To obtain the target user's node representation for the next layer $l+1$, $\tilde{h}_u^l$ is aggregated with learned weights at current layer l using a non linear transformation,

$$h_u^{(l+1)} = ReLU(\mathbf{W}^{(l)} \tilde{h}_u^l) \tag{10}$$

We stack each node L times to obtain the final target user node representation, $h_u^{(L)}$. This design preserves volunteer's individual preference h_u while harnessing the impact of each neighbour h_j.

The dynamic feature graph G is constructed such that the nodes correspond to the volunteer and her neighbours, that is, for target user u, with $\mid N_K \mid$ neighbours, the graph consists of $\mid N_k \mid +1$ nodes. The attention mechanism is used to determine the weight of neighbour k on target volunteer u that is propagated along the edges of the network. The initial representation of target volunteer h_n is used as node features for layer 0, denoted as $h_u^{(0)}$. Similarly, a neighbour's node features $h_k^{(0)}$ are given by the combined representation of the short and long-term interests $S^{u,k}$. While the features of the target volunteer changes whenever the volunteer participates in a new task, the neighbour's features remain unchanged for the duration of the timestep $(T+1)$. After computing weights for all layers, the final representation of a volunteer is obtained by combining the volunteer network and the volunteer's most recent task participation using a fully connected network. Formally:

$$\hat{h}_n = D_1\left[h_n; h_u^{(L)}\right] \tag{11}$$

D_1 is a linear transformation matrix and $\hat{h}_n$ is the final representation of the volunteer's current preference. Softmax is used to compute the probability of the next organizer for a volunteer and provide recommendations.

4 Experiments and results

4.1 Dataset

We use data collected from "Anti-Pandemic Pioneers" project (a.k.a. Pioneers) (Zhang et al., 2022), a mobile-based platform used for mobilizing volunteers and recording their participation activities in Shenzhen, China. The platform was initially launched to coordinate volunteer efforts at the start of Covid-19 epidemic in February 2020, and it has since been expanded to include other volunteer tasks. On Pioneers, users can organize short-term group activities and perform micro-tasks within small communities. Moreover, Pioneers is a decentralized platform such that each user can either select an existing task or list a new task that other users can participate in. We categorically treat the former as the volunteer and the latter as the organizer, even though these roles can be used interchangeably, that is, a volunteer can also be an organizer, and vice-versa.

When computing session-based recommendations, optimizing session segmentation is an important task that has been overlooked in previous studies. For example, Song et al. (2019) segmented the data into sessions based on empirical frequency of users' consumption of items. On the contrary, we used the Shannon Entropy (Shannon, 1948) to characterize the randomness and uncertainty of session data. We define session segmentation, $\mathcal{S}$ as follows:

$$\mathcal{S} = \frac{1}{\mid \mathcal{D} \mid} \sum_{i=1}^{N} H(X_i) \tag{12}$$

$$H(X_i) = -\sum_{x \in \mathcal{X}} p(x) \log(x) \forall i \in \mathcal{D} \tag{13}$$

Thus, to determine the most suitable time window, we computed the entropy for $\mathcal{D} =$ [1,7,14,21,28]-day segments and selected session segmentation corresponding to average entropy. Our goal was to create sessions that are not too short and easy to predict but also not too long and complex for the network to train. As a result, we segmented the data based on 7 day long sessions ordered by timestamps, from February 2020 to January 2021. Moreover, sessions of length 1 were filtered from the data (Hidasi et al., 2016). To avoid outliers, we set maximum session length empirically to 30. Table 1 shows the descriptive statistics for Pioneers data.

Table 1 Summary Statistics for Pioneers Data Before Processing

Description	Stats
Total users	80 043
# of Total organisers	7 391
Total # of Events	2 539 609
Total # users in volunteer network	72 141
Total Edges	800 430
Average # of friends	10
Average # of events per user	31.73

4.2 Parameter settings

The model was implemented in TensorFlow, a popular python library for machine learning. To maintain the order of sessions in training and validation, we reserved the last d days for validation. We set d to 14 and randomly separated the validation sessions into two halves. The first half was used as validation set to tune the hyper-parameters and the second half as testing set. We used Adam for optimization due to its effectiveness (Kingma and Ba, 2014), and set the learning rate to 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and ϵ=1e-07 as suggested in Tensorflow. The decay rate was set to 0.98 for every 400 steps. The sample size for the neighbourhood was set to 10 and 5 in the first and second convolutional layers respectively. Furthermore, the LSTM hidden units were empirically set to 100, and the dimensions of the volunteer and organiser for the model were also set to 100, following Hidasi et al. (2016). We use ReLU as the activation function and to avoid overfitting, a dropout rate of 0.2 was used. The early stopping strategy was adopted to end training when there is no longer improvement in the validation accuracy. Training consists of 10 epochs with 1 500 iterations in each epoch. We conducted 5-fold cross validation on all data to obtain the overall accuracy and standard error of our proposed method. After preprocessing, Table 2 shows the final data used for training, validation and testing.

This study uses recall to evaluate the performance of the model. Recall measures the proportion of relevant items found in the top K recommendations. We set $K = 3$, 5 and 10. We compare the results against four state of the art recommendation baselines.

- **GRU4Rec** is a sequence-based method proposed by Hidasi et al. (2016). The algorithm takes the actual state of a session as input and uses a Gated Recurrent Unit (GRU) to update the hidden state and output vector. In our volunteer data, we construct sessions by segmenting each volunteer's total task participation history using a one week rolling window. To maintain consistency, we also remove sessions of length 1 from the training and test sets. We leave the last week for testing.

Table 2 Summary Statistics for Pioneers Data After Processing

Description	Training	Validation	Testing
Events	1 286 051	53 586	54 523
Sessions	350 096	14 274	14 417
Organisers	2 354	864	876
Avg. length	3.67	3.75	3.78

- **ItemKNN** is a collaborative filtering technique that analyzes user-item matrix to discover the relationships between items, and indirectly compute recommendations based on these relationships (Sarwar et al., 2001). We use cosine similarity between vectors of organizers, and recommend organizers to volunteers in the setting "volunteers who participated under this organizer also participated under these organizers".
- **BPR-MF** proposed by Rendle et al. (2012) is a matrix factorization method that uses pairwise ranking loss via stochastic gradient descent. Since matrix factorization can not be used directly for session-based recommendation, we improvise by letting a new session be represented by average latent factors of items that have occurred in the session so far. We then compute the recommendation score as the average of similarities between latent factors of a candidate item and the items in the session so far.
- **POP** recommends the most popular items in the training set (Li et al., 2017; Hidasi et al., 2016). We set the number of popular items to 100. Despite its simplicity, this approach often serves as a robust baseline in specific domains.

4.3 Results and discussion

The results in Table 3 shows the best performing models. We compare the results when all data is used against break point intervals implemented on Futian district. Details on motivation and construction of structural breaks are in Subsection 4.4. Our proposed method demonstrates substantial gain over baselines across all data sets except in BP 2. Our analysis show that VolRec outperforms the second-best model by approximately 27% on all data. Furthermore, using Recall@10 the 5-fold cross validation results have a standard error of 0.00095. The low standard error indicates the robustness of our method and underscores the reliability and performance consistency of VolRec. The results suggests that the network architecture's ability to incorporate neighbourhood information allows it to propagate important information that improves recommendation accuracy and ultimately outperform the baselines. More importantly, using temporal dynamics to model sequential user interactions as session embeddings helps to express the changing structure of volunteer participation. GRU4Rec performs better than other baselines both on district data and when all data is used. This indicates that GRU4Rec is a strong baseline on longer sessions.

When compared against break point intervals which have shorter sessions, BPR-MF is very competitive and achieves the best performance in BP 2. It is interesting to note that increasing the value of K seems to benefit BPR-MF and POP the most, in both long and short sessions. ItemKNN shows good performance when all data is used but the accuracy gradually declines in small samples. This is expected since ItemKNN does not consider temporal information and only depends on item similarity to make recommendations. When sessions are shorter or sparse, GRU4Rec performs poorly than most baselines since it only relies on volunteer's own sessions to learn the embeddings and make recommendations. However, our proposed method overcomes this challenge by utilising neighbourhood information in the volunteer network. The strategy to construct a probabilistic volunteer network in the absence of a predefined social network data is crucial in capturing volunteers' implicit relationships and subsequent dynamic preferences. Additionally, the use of an attention mechanism to learn relative weights between connected nodes underscores the significance of neighbour's influence for each target volunteer.

Table 3 We show Recall@K for baseline models and VolRec on different data sets

	All data			Futian District Data		
Model	Recall@3	Recall@5	Recall@10	Recall@3	Recall@5	Recall@10
GRU4Rec	0.77195	0.7767	0.7852	0.6702	0.6781	0.6784
Item-KNN	0.6702	0.6753	0.6813	0.5028	0.5142	0.5206
BPR-MF	0.2848	0.5127	0.7348	0.1699	0.3633	0.5693
POP	0.1951	0.2618	0.3451	0.0690	0.1166	0.1458
VolRec	**0.9835**	**0.9931**	**0.9946**	**0.9666**	**0.9815**	**0.9851**
	BP 1			BP 2		
GRU4Rec	0.3601	0.6865	0.8574	0.1117	0.1117	0.1210
Item-KNN	0.5021	0.5037	0.5108	0.5041	0.5110	0.5151
BPR-MF	0.3119	0.6013	0.6506	0.5493	0.9628	**0.9926**
POP	0.3028	0.3773	0.5295	0.4283	0.4544	0.5624
VolRec	**0.9286**	**0.9591**	**0.9916**	**0.9566**	**0.9916**	0.9923
	BP 3			BP 4		
GRU4Rec	0.1171	0.1518	0.2021	0.3336	0.3449	0.3569
Item-KNN	0.5058	0.5163	0.5218	0.5088	0.5166	0.5288
BPR-MF	0.2458	0.4466	0.6469	0.1894	0.2628	0.3838
POP	0.1131	0.1825	0.2240	0.0618	0.1019	0.1604
VolRec	**0.9214**	**0.9264**	**0.9207**	**0.8963**	**0.9790**	**0.9781**
	BP 5					
GRU4Rec	0.1171	0.2690	0.2821			
Item-KNN	0.5044	0.5144	0.5223			
BPR-MF	0.1274	0.2899	0.4150			
POP	0.0705	0.1930	0.2492			
VolRec	**0.9401**	**0.9790**	**0.9781**			

We also compare the organizer-volunteer recommendation performance of the models on 5 break point intervals discovered on Futian District Data, `BP 1-5`. Our proposed **VolRec** model consistently outperforms the baselines

Bold denotes the best result and underline denotes the second best result

4.4 Change-point detection

To gather more insights and justify the architecture of our proposed method, we study the performance of this approach under different settings. We believe that the behavior of volunteers differed across task labels. For instance, since the data predominantly covers Covid-19 pandemic period, preliminary analysis shows that major policy changes were implemented by the government which impacted volunteer participation behavior. We observed peaks and gradual declines in volunteer task participation that corresponded to the time of major policy announcements.

Meanwhile, the number of daily volunteers' task participation could be indicative of the shift in focus and priorities towards different task categories. Therefore, it is essential to examine whether the distribution of task types impacted the overall performance of the model. To achieve this, we first learn topics from task descriptions using Latent Dirichlet Allocation (Blei et al., 2003). We discovered six distinctive task labels namely (1) public transportation, (2) community volunteering, (3) business reopening, (4) public health education, (5) envi-

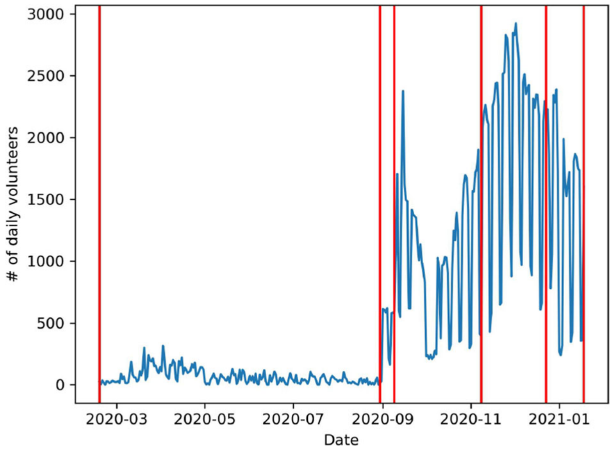

Fig. 6 We show that the behaviour of volunteers and organizers varied significantly across time. We discovered 4 structural break dates, separating data into 5 break point intervals

ronmental protection, and (6) Covid-19 response. We eliminated smaller topics, (task label 0) that could not be identified with these task labels.

In the second step, we chose to focus on Futian district in Shenzhen. Our choice to evaluate the model at district level is motivated by McPherson et al. (2001) who showed that volunteers in the same region tend to have similar behavioral patterns. Moreover, policies were also implemented at district level, which prompts us to consider volunteer participation at that level.

Third, we use volunteers daily data to analyze their task preferences and participation behaviour. We also tried daily organizer data and other signals and the results were similar. Given the data in Fig. 6, it is clear that certain periods experienced high and active volunteer participation while others did not. For example, the total number of daily volunteers barely reached 500 a day since February 2021, but spiked to over 2 000 in September 2021. We observed a similar pattern in organizer task issuance data. This phenomenon is closely related to Covid-19 policy changes at government level, which in turn impacted volunteer participation behaviour.

As a result, we study these different periods separately. Precisely, we implement change point detection (Truong et al., 2020), a popular technique in signal processing to detect structural breaks in data. Four break dates were identified which reflect periods of significant deviation in volunteer behaviour pattern as shown in Fig. 6. In Figs. 7 - 11 we present the break point intervals separately and their respective distribution of task labels.

4.5 Impact of volunteers task preference

In the first part of our exploration, we trained the model using all task labels in each break point interval. Next, we compute the inference for each break point interval and the results are reported in Table 3. It is interesting to note that using our method, on average, the second break point interval (BP 2) achieves the highest recall while interval four (BP 4) has the lowest score.

We note that since the behavior of volunteers and organizers varied significantly across time (Fig. 6), it also reflects changes in priority towards certain task types within a given

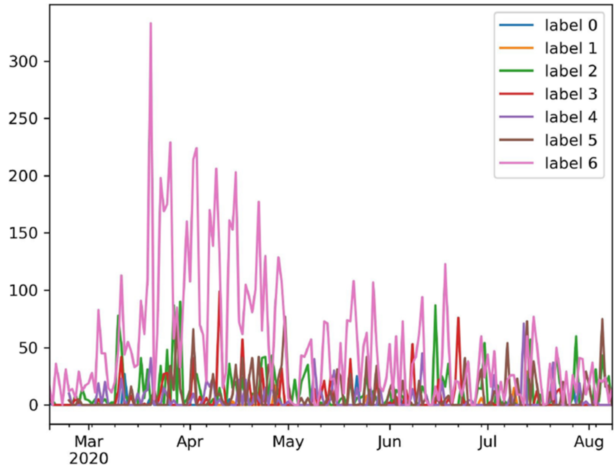

Fig. 7 Distribution of task labels for break point interval 1

interval as shown in Fig. 7-11. Previous studies such as Hidasi et al. (2016) fail to address temporal changes in user preferences which is informative when analyzing model's performance. To gain further insights, we hypothesize that volunteer task preference, τ, affects the model's accuracy in each break point interval. We separate the testing data for each sample based on task labels to determine their impact on model performance. The results are quite striking. Figure 12 reports the performance of task types τ, and their corresponding break point intervals. We observe that task types 3 and 1 have the highest recall while 2 and 5 display the worst performance.

We further analyze the level of uncertainty embedded in daily volunteer task participation by evaluating both the entropy and standard deviation of each task type within a break point interval. The results of entropy and standard deviation exhibit a similar pattern, and we only report results of entropy in Fig. 13 .

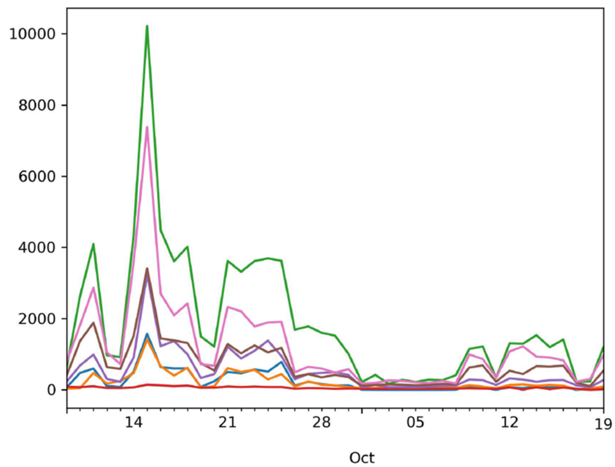

Fig. 8 Distribution of task labels for break point interval 2

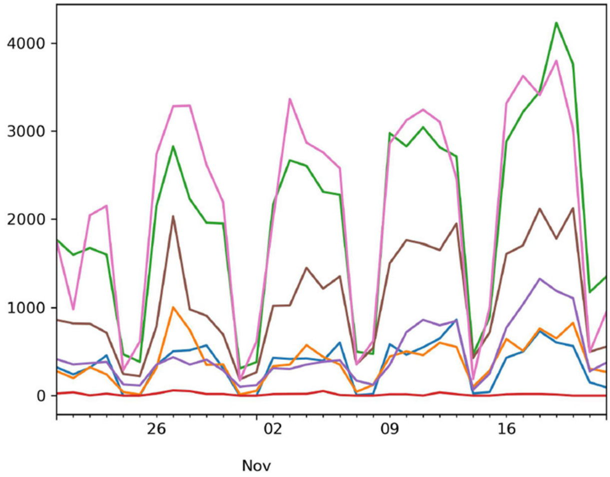

Fig. 9 Distribution of task labels for break point interval 3

We observe that where the accuracy is high, the entropy is generally low and vice-versa. This inverse relationship is a common phenomenon in machine learning, where it is difficult to learn and predict data with high randomness. In our case, the uncertainty in the structure of daily volunteer participation increased the level of difficulty in getting accurate predicted outcomes. We also analyzed task issuance by organizers and arrived at the same conclusion. These results demonstrate the strong impact of task preference, as indicative of volunteer's dynamic behaviour in determining accuracy. This study is the first to explore in detail the relationship between changes in volunteer's task preference and prediction accuracy.

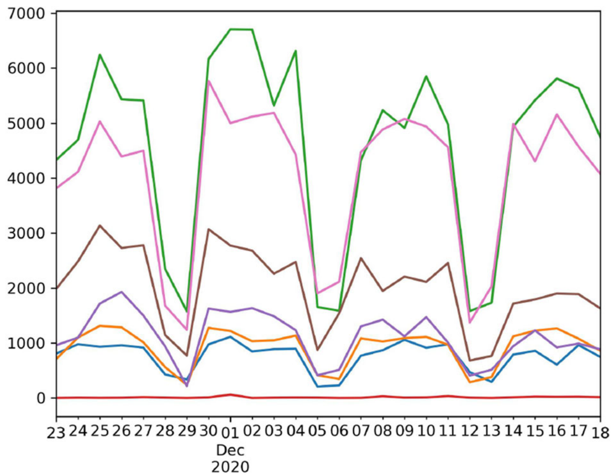

Fig. 10 Distribution of task labels for break point interval 4

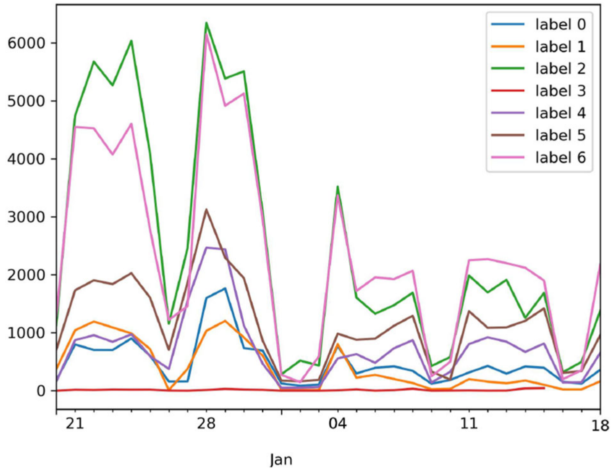

Fig. 11 Distribution of task labels for break point interval 5

4.6 Volunteer group recommendation

VolRec is a comprehensive session-based model which can be adapted to predict group participation for volunteers. Modeling volunteer group participation is essential in understanding volunteer communities and their dynamic structures. For instance, if a volunteer i participates in a task with n other volunteers, our objective is to determine the most likely volunteers that i will co-participate with in the next sessions. We construct volunteer-to-volunteer network graph and utilize temporal dynamics to create sessions. For this analysis we empirically set session length to 1 day and leave the last 10 days for testing. We conduct experiments on the break points discovered in Subsection 4.4 and the results are shown in Table 4 below.

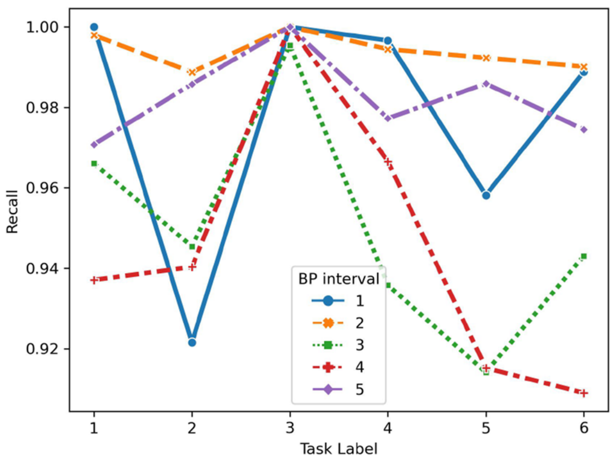

Fig. 12 Visualization of task label performance for different break points. Results show that the performance of task types, τ, differed significantly across break point intervals

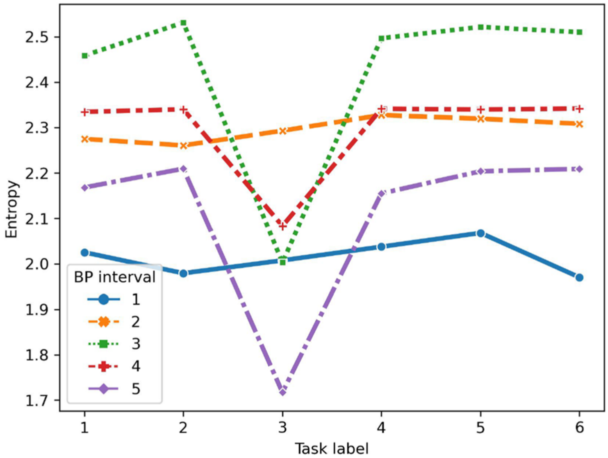

Fig. 13 The entropy of volunteers' daily task participation for each break point interval

Overall, group volunteer recommendation accuracy is lower than organizer-volunteer recommendation's. This is expected since the number of organizers is less than that of volunteers and consequently more difficult to predict as the number of predicted outcomes is large. However, break point intervals 2 and 4 have the highest recall while 1 and 3 achieve the lowest. Our model is able to learn structural relationships of group participants and recommend volunteers to co-participate with in the next sessions with reasonable accuracy.

Moreover, we tested VolRec on volunteer-task type prediction and we achieved a perfect recall on all break point intervals. Besides, VolReck can be adapted to model and predict communities for volunteers.

4.7 Ablation study

4.7.1 Individual vs neighbourhood influence

In the second phase of our exploration, we conduct an ablation study to justify the influence of neighbourhood on the recommendations. In our previous experiments, we fixed the volunteer neighbourhood to 10 (refer Subsection 3.3). We now use individual features only without neighbourhood information. The model without neighbourhood information is similar to GRU4Rec and the results are reproduced in Table 3. When we use neighbourhood information, the performance significantly improves from 0.785 to 0.989 using Recall@10 when all data is used. These results express the importance of utilizing neighbourhood information to obtain good recommendations in online volunteer platforms. It is therefore important to model individual and neighbours' interests to achieve higher recommendation performance.

Table 4 Results of Volunteers' Group Recommendation. We show that **VolRec** is able to learn implicit group relationships and make recommendations with reasonable accuracy

Break points	BP 1	BP 2	BP 3	BP 4	BP 5
Recall@10	0.4828	0.5599	0.5105	0.5644	0.5200

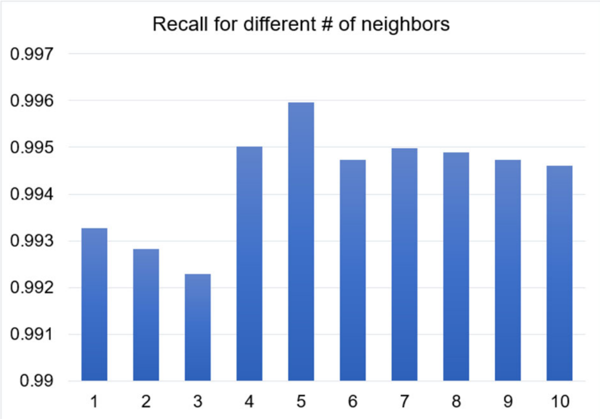

Fig. 14 The impact of the number of neighbors on model performance. We show experimentally that the optimal number of neighbours is 5

4.7.2 Influence of neighbours

The context of our problem is slightly different from the general online recommendations in that our online volunteer data does not come with pre-defined neighbours. As noted earlier, we constructed volunteer network from the existing volunteer data that we already had, (refer Subsection 3.1). To understand the influence of neighbours in detail, it is important to determine the optimal number of neighbours for each volunteer, which in turn improves efficiency of the algorithm. We tested the model by incrementing the number of neighbours from 1 to 10. In Fig. 14, we compare the recommendation performance for a given number of neighbours.

Results clearly show that the optimal number of neighbours for our volunteer network is 5. While adding neighbourhood information increases recommendation performance, further analysis indicates that continuously adding neighbours after 1 does not significantly improve the results. This is expected since our construction of the volunteer network is based on co-participation, a feature that is shared by all other neighbours. Therefore, increasing the number of neighbours does not necessarily bring new valuable information. We also noted slight improvement in the efficiency of the algorithm when only one neighbour node is used for each target volunteer.

5 Conclusions and future work

In this paper we proposed VolRec, a session-based recommendation for large online volunteer networks. We found that volunteer's dynamic interests can be represented as a combination of their most recent sessions and the long and short-term sessions of their neighbours. We apply graph neural networks and propagate the weights using an attention mechanism. We show that VolRec outperforms several baselines in recommending organizers to volunteers due to its ability to capture dynamic users and neighbours' interests. We investigate the performance of VolRec on break point intervals that represent structural changes in volunteer behavior. The performance of our model varied across different break points. To gain further insight, we evaluated VolRec on task types discovered in each break point interval separately. We found that the distribution of task types in the intervals impacted the model's performance. Additionally, we characterised the uncertainty in volunteers' daily task participation using entropy and discovered an inverse relationship between task type performance and entropy. This phenomenon is a common paradox in machine learning where it is difficult to predict data with a high degree of randomness. We conducted an ablation study to determine the influence of neighbours and the optimal number of neighbours. Results demonstrate that including neighbourhood information increased the recommendation accuracy by $\sim 26\%$. While the optimal number of neighbours is 5, we note that increasing the neighbours after 1 did not significantly improve the results since the construction of our volunteer network is based on co-participation, a feature that is shared by all other neighbours. Furthermore, we conducted experiments to show that VolRec can be adapted to recommend volunteers to task types and communities. In future research, it would be interesting to expand the application of our proposed method to various other online platforms, such as Wikipedia or Stack Overflow. This would allow us to recommend the next articles to edit or questions to answer, and conduct further analyses to gain deeper insights into the complex latent relationships in online data. In other future direction, we aim to address the lower performance observed when the entropy of task types is high.

Acknowledgements We thank the Shenzhen Municipal Organizational Department for providing the volunteer data.

Funding This study is supported in part by the Tsinghua SIGS Scientific Research Start-up Fund (Grant Grant QD2021012C) and Natural Science Foundation of China (Grant 62001266)

Data Availability The datasets generated during and/or analysed during the current study are available in the Harvard Dataverse repository, https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/YUOOBB

Code Availability The code and sample data used in this paper are available on our Github repository, https://github.com/TauraiUCB/VolRec

Declaration

Conflict of interest/Competing interests The authors declare no competing interests

Ethics approval Not applicable as this study did not involve human participants

Consent to participate Not applicable as this study did not involve human participants

Consent for publication The authors grant the Publisher an exclusive licence to publish the article

References

An, M., Wu, F., Wu, C., et al. (2019). Neural news recommendation with long-and short-term user representations. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 336–345. https://doi.org/10.18653/v1/P19-1033

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 993–1022.

Choi, M., Kim, J., Lee, J., et al. (2022). S-Walk: accurate and scalable session-based recommendation with random walks. *In: Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining,* pp. 150–160. https://doi.org/10.1145/3488560.3498464

Donkers, T., Loepp, B., & Ziegler, J. (2017). Sequential user-based recurrent neural network recommendations. *In: Proceedings of the eleventh ACM conference on recommender systems,* pp. 152–160. https://doi.org/10.1145/3109859.3109877

Eagle, N., Pentland, A., & Lazer, D. (2008). Mobile phone data for inferring social network structure. *In: Social Computing, Behavioral Modeling, and Prediction,* pp. 79–88. https://doi.org/10.1007/978-0-387-77672-9_10

Fan, W., Ma, Y., Li, Q., et al. (2019). Graph neural networks for social recommendation. *In The World Wide Web Conference,* pp. 417–426. https://doi.org/10.1145/3308558.3313488

Gao, C., Zheng, Y., Li, N., et al. (2023). A survey of graph neural networks for recommender systems: Challenges, methods, and directions. *ACM Transactions on Recommender Systems, 1*(1), 1–51. https://doi.org/10.1145/3568022

Gao, C., Wang, X., He, X., et al. (2022). Graph neural networks for recommender system. *In: Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining,* pp. 1623–1625. https://doi.org/10.1145/3488560.3501396

Guo, L., Yin, H., Wang, Q., et al. (2019). Streaming session-based recommendation. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 1569–1577. https://doi.org/10.1145/3292500.3330839

Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30

He, X., Liao, L., Zhang, H., et al. (2017). Neural collaborative filtering. In: Proceedings of the 26th International Conference on World Wide Web, pp. 173–182. https://doi.org/10.1145/3038912.3052569

He, Y., Mao, Y., Xie, X., et al. (2022). An improved recommendation based on graph convolutional network. *Journal of Intelligent Information Systems, 59*(3), 801–823. https://doi.org/10.1007/s10844-022-00727-3

Hidasi, B., Karatzoglou, A., Baltrunas, L., et al. (2015). Session-based recommendations with recurrent neural networks. arXiv:1511.06939

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv:1412.6980

Kipf, T. N., & Welling, M. (2016). Semi-supervised classification with graph convolutional networks. arXiv:1609.02907

Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer, 42*(8), 30–37. https://doi.org/10.1109/MC.2009.263

Lei, J., Li, Y., Yang, S., et al. (2022). Two-stage sequential recommendation for side information fusion and long-term and short-term preferences modeling. *Journal of Intelligent Information Systems, 59*, 657–677. https://doi.org/10.1007/s10844-022-00723-7

Li, J., Ren, P., Chen, Z., et al. (2017). Neural attentive session-based recommendation. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, pp. 1419–1428. https://doi.org/10.1145/3132847.3132926

Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing, 7*(1), 76–80. https://doi.org/10.1109/MIC.2003.1167344

Liu, J., Dolan, P., & Pedersen, E. R. (2010). Personalized news recommendation based on click behavior. In: Proceedings of the 15th International Conference on Intelligent User Interfaces, pp. 31–40. https://doi.org/10.1145/1719970.1719976

Liu, Q., Zeng, Y., Mokhosi, R., et al. (2018). STAMP: short-term attention/memory priority model for session-based recommendation. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 1831–1839. https://doi.org/10.1145/3219819.3219950

Locke, M., Ellis, A., & Smith, J. D. (2003). Hold on to what you've got: the volunteer retention literature. *Voluntary Action, 5*(3), 81–99.

Marston, C., Renedo, A., & Miles, S. (2020). Community participation is crucial in a pandemic. *The Lancet, 395*(10238), 1676–1678. https://doi.org/10.1016/S0140-6736(20)31054-0

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology, 27*(1), 415–444. https://doi.org/10.1146/annurev.soc.27.1.415

Rendle, S., Freudenthaler, C., & Schmidt-Thieme, L. (2010). Factorizing personalized markov chains for next-basket recommendation. In: Proceedings of the 19th International Conference on World Wide Web, pp. 811–820. https://doi.org/10.1145/1772690.1772773

Rendle, S., Freudenthaler, C., Gantner, Z., et al. (2012). BPR: Bayesian personalized ranking from implicit feedback. arXiv:1205.2618

Sarwar, B., Karypis, G., Konstan, J., et al. (2001). Item-based collaborative filtering recommendation algorithms. In: Proceedings of the 10th International Conference on World Wide Web, pp. 285–295. https://doi.org/10.1145/371920.372071

Shani, G., Heckerman, D., Brafman, R. I., et al. (2005). An MDP-based recommender system. *Journal of Machine Learning Research*, 6(9)

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal, 27*(3), 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Song, C., Wang, B., Jiang, Q., et al. (2021). Social recommendation with implicit social influence. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1788–1792. https://doi.org/10.1145/3404835.3463043

Song, W., Xiao, Z., Wang, Y., et al. (2019). Session-based social recommendation via dynamic graph attention networks. *In: Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining,* pp. 555–563. https://doi.org/10.1145/3289600.3290989

Truong, C., Oudre, L., & Vayatis, N. (2020). Selective review of offline change point detection methods. *Signal Processing, 167*, 107299. https://doi.org/10.1016/j.sigpro.2019.107299

Veličković, P., Cucurull, G., Casanova, A., et al. (2017). Graph attention networks. arXiv:1710.10903

Wang, S., Zhang, Q., Hu, L., et al. (2022). Session-based Recommendations: Challenges, Approaches, Applications and Opportunities. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 3425–3428. https://doi.org/10.1145/3477495.3532685

Wang, Z., Liao, J., Cao, Q., et al. (2014). Friendbook: a semantic-based friend recommendation system for social networks. *IEEE Transactions on Mobile Computing, 14*(3), 538–551. https://doi.org/10.1109/TMC.2014.2322373

Wu, S., Tang, Y., Zhu, Y., et al. (2019). Session-based recommendation with graph neural networks. *In Proceedings of the AAAI Conference on Artificial Intelligence, 33*(1), 346–353. https://doi.org/10.1609/aaai.v33i01.3301346

Xu, K., Li, C., Tian, Y., et al. (2018). Representation learning on graphs with jumping knowledge networks. In: International Conference on Machine Learning, pp. 5453–5462.

Zhang, A., Zhang, K., Li, W., et al. (2022). Optimising self-organised volunteer efforts in response to the COVID-19 pandemic. Humanities and Social Sciences Communications, 9(1). https://doi.org/10.1057/s41599-022-01127-2